



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/178589/>

Version: Accepted Version

Proceedings Paper:

Cooper, C, Dyer, M and Greenhill, C (2021) A Triangle Process on Regular Graphs. In: Combinatorial Algorithms. 32nd International Workshop, IWOCA 2021, 05-07 Jul 2021, Ottawa, ON, Canada. , pp. 310-323. ISBN: 978-3-030-79986-1. ISSN: 0302-9743. EISSN: 1611-3349.

https://doi.org/10.1007/978-3-030-79987-8_22

© Springer Nature Switzerland AG 2021. This is an author produced version of an article, published in Combinatorial Algorithms. IWOCA 2021. Lecture Notes in Computer Science. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

A triangle process on regular graphs^{*}

Colin Cooper¹, Martin Dyer², and Catherine Greenhill³

¹ Department of Informatics, King's College, London WC2R 2LS, U.K.
`colin.cooper@kcl.ac.uk`

² School of Computing, University of Leeds, Leeds LS2 9JT, U.K.
`m.e.dyer@leeds.ac.uk`

³ School of Mathematics and Statistics, UNSW Sydney, NSW 2052, Australia.
`c.greenhill@unsw.edu.au`

Abstract. Switches are operations which make local changes to the edges of a graph, usually with the aim of preserving the vertex degrees. We study a restricted set of switches, called triangle switches. Each triangle switch creates or deletes at least one triangle. Triangle switches can be used to define Markov chains which generate graphs with a given degree sequence and with many more triangles (3-cycles) than is typical in a uniformly random graph with the same degrees. We show that the set of triangle switches connects the set of all d -regular graphs on n vertices, for all $d \geq 3$. Hence, any Markov chain which assigns positive probability to all triangle switches is irreducible on these graphs. We also investigate this question for 2-regular graphs.

Keywords: Regular graphs · Triangles · Markov chains · Irreducibility.

1 Introduction

Generating graphs at random from given classes and distributions has been the subject of considerable research. See, for example, [1,2,3,5,8,10,11,15,16,18,17,21]. Generation using Markov chains has been a topic of specific interest in this context, in particular Markov chains based on switches of various types, for example [2,5,8,10,15,16,18,17,21]. Switches delete a pair of edges from the graph and insert a different pair on the same four vertices. They have the important property that they preserve the degree sequence of the graph. Thus they are useful for generating regular graphs, or other graphs with a given degree sequence. Markov chains also give a dynamic reconfigurability property, which is useful in applications, for example [5,10,17]. For any such Markov chain, two questions arise. First, can it generate any graph in the chosen class? (Formally: is the Markov chain irreducible?) Secondly, we might wish to estimate its rate of convergence to the chosen distribution. (Formally: what is the mixing time of the chain?)

In the applied field of social networks, the existence of triangles (3-cycles) is seen as an indicator of mutual friendships [12,14]. Such networks can alter

^{*} Dyer supported by EPSRC grant EP/S016562/1, “Sampling in hereditary classes”, Greenhill by Australian Research Council grant DP190100977.

in a dynamic fashion as pairs of vertices friend or unfriend each other based on a mutual acquaintance. However, many random graph models, or processes for producing random graphs, tend to produce graphs with few triangles. This is true for any process which generates sparse graphs with a given degree sequence (approximately) uniformly at random. For example, the expected number of triangles is constant for d -regular graphs, when d is constant [4]. In this paper, we study a restricted set of switches, called *triangle switches*, and consider any reversible Markov chain whose transitions are exactly the triangle switches. Triangle switches were introduced in [6] in the context of cubic graphs, and examples were given of Markov chains using triangle switches with transition probabilities assigned to encourage the formation of triangles. It was proved in [6, Section 4] that it is possible to generate cubic graphs using this approach which have $\Omega(n)$ triangles in $O(n)$ steps of the Markov chain.

In this paper we address the first question posed above (“is the Markov chain irreducible?”) for such chains on the state space of d -regular graphs, for any d . Note that the answer to this question is independent of the probabilities assigned to each triangle switch by the Markov chain, as it is a property of the undirected graph underlying the Markov chain. We leave the mixing question for future research, noting only that tight bounds on mixing time seem hard to come by in this setting. The recent paper [21] is a notable exception.

The proofs in [6] do not easily generalise to regular graphs of arbitrary d , though the main approach in our proof of irreducibility comes from [6]. If a component of a d -regular graph is a clique (that is, a complete subgraph) then it must be isomorphic to K_{d+1} . We call such a component a *clique component*. Our approach is to show that starting from an arbitrary d -regular graph, triangle switches can be used to increase the number of clique components. Furthermore, we show how to alter the set of vertices in a given clique component using triangle switches. After creating as many clique components as possible, there is at most one additional component C , which must satisfy $d + 1 < |C| < 2(d + 1)$. We call such a component a *fragment*. We prove that triangle switches connect the set of all fragments on a given vertex set. In the cubic case, this last step is simpler as the only possible fragments are $K_{3,3}$ and \overline{C}_6 .

Our result can be viewed as solving a particular *reconfiguration* problem for regular graphs. Reconfiguration is a topic of growing interest in discrete mathematics. For an introduction to the topic, and a survey of results, see [19]. We note that reconfiguration problems can be as hard as PSPACE-complete, in general. Our results show that there is a polynomial time algorithm to construct a path of triangle switches between any two d -regular graphs on n vertices.

The plan of the paper is as follows. In Section 1.1, we define and review switches and restricted switches, in particular triangle switches, and state our main result, Theorem 1. For most of the paper we assume that $d \geq 3$. In Section 2 we show that the set of all fragments with a given vertex set is connected under triangle switches. In Section 3, we show that triangle switches can be used to create a clique component, starting from any d -regular graph with at least $2(d+1)$ vertices. In Section 4 we show how to relabel the vertices in clique components

using triangle switches, and hence complete our proof of irreducibility. Finally, in Section 5, we consider the irreducibility question for d -regular graphs with $d \leq 2$. Many proofs are sketched or omitted here. For full details see [7].

Definitions, notation and terminology. The notation $[k]$ will denote the set $\{1, 2, \dots, k\}$, for any integer k . Given a set V of vertices, let $V^{\{2\}}$ be the set of unordered pairs of distinct elements from V . A graph $G = (V, E)$ on vertex set $V(G)$ has edge set $E(G) \subseteq V^{\{2\}}$. We usually denote $|V|$ by n . We use the notation xy as a shorthand for the unordered pair $\{x, y\}$, whether or not this pair is an edge. If $E' \subseteq E$ and $V' = \{v \in V : v \in e \in E'\}$, then $G' = (V', E')$ is a *subgraph* of G . Given any vertex subset $U \subseteq V$, the subgraph $G[U]$ induced by U has vertex set U and edge set $E' = U^{\{2\}} \cap E$. If $|U| = k$ and $G[U]$ is a k -cycle, then we say that $G[U]$ is an *induced C_k* .

We will write $G \cong H$ to indicate that graphs G and H are isomorphic. Given a graph $G = (V, E)$, the *complement* of G is the graph $\overline{G} = (V, \overline{E})$ with $\overline{E} = V^{\{2\}} \setminus E$. An edge of \overline{G} will be called a *non-edge* of G .

The *distance* $\text{dist}(u, v)$ between two vertices u and v is the number of edges in a shortest path from u to v in G , with $\text{dist}(u, v) := \infty$ if no such path exists. The maximum distance between two vertices in G is the *diameter* of G , and G is connected if it has finite diameter. The *component* C of G containing v is the largest connected induced subgraph of G which contains v .

Given a graph $G = (V, E)$ and vertex $v \in V$, let $N_G(v) = \{u : uv \in E\}$ denote the neighbourhood of v , and let $\deg_G(v) = |N_G(v)|$ denote the degree of v in G . The *closed neighbourhood* of v is $N_G[v] := N_G(v) \cup \{v\}$. We sometimes drop the subscript and write $N(v)$ or $N[v]$.

Say that G is *regular* if every vertex has the same degree, and if $\deg_G(v) = d$ for all $v \in V$ then we say that G is d -regular. As already stated, $\mathcal{G}_{n,d}$ will be the set of all d -regular graphs with vertex set $V = [n]$. Note that $\mathcal{G}_{n,d}$ is non-empty if and only if either d or n is even. This result seems to be folklore, but is easy to prove. Necessity is implied by edge counting, and sufficiency by a direct construction. An indirect proof can be found in [22, Prop. 1]. As usual, $K_{d+1} \in \mathcal{G}_{d+1,d}$ denotes the complete graph on $d+1$ vertices, and $K_{d,d} \in \mathcal{G}_{2d,d}$ denotes the complete bipartite graph on $d+d$ vertices. A graph in $\mathcal{G}_{n,d}$ with $d+1 < n < 2(d+1)$ will be called a *fragment*. Note that K_{d+1} is not a fragment.

We often regard a graph $G \in \mathcal{G}_{n,d}$ as *layered*, in the following way. Let v be a given (fixed) vertex of a d -regular graph $G = (V, E)$, where $n = |V| \geq 2(d+1)$, and let $C \subseteq V$ determine the component $G[C]$ of G such that $v \in C$. We regard C as partitioned by distance from v , with $V_i = \{u \in C : \text{dist}(v, u) = i\}$. Thus $V_0 = \{v\}$ and $V_1 = N(v)$, so $|V_0| = 1$ and $|V_1| = d$. Since V_2 appears frequently in the proof, we will denote $|V_2|$ by ℓ . By definition, C is a disjoint union $C = \bigcup_{i \geq 0} V_i$, and $|C| = \sum_{i \geq 0} |V_i|$. Let $G_i = G[V_i] = (V_i, E_i)$, and note that $G_0 = (\{v\}, \emptyset)$. Let $N'(u)$ be the neighbourhood of $u \in V_i$ in G_i , i.e. $N'(u) = N(u) \cap V_i$, and let $d'(u) = |N'(u)|$ be the degree of u in G_i . We omit explicit reference to i in this notation, since it is implicit from $u \in V_i$. Given $u \in V_i$, we denote the set of non-neighbours of u in G_i by $\overline{N}'_i(u)$.

We will regard the edges of $G[C]$ from V_i to V_{i+1} as being directed away from the designated vertex v . Under this convention, for $u \in V_i$, $\text{In}(u)$ is the neighbour set of u in V_{i-1} , and $\text{Out}(u)$ is the neighbour set of u in V_{i+1} . Thus, if $u \in V_i$, $\text{In}(u) = \text{N}(u) \cap V_{i-1}$ and $\text{Out}(u) = \text{N}(u) \cap V_{i+1}$. Then let $\text{id}(u) = |\text{In}(u)|$ be the in-degree of $u \in V_i$ and $\text{od}(u) = |\text{Out}(u)|$ the out-degree of u . Thus $\text{N}(u) = \text{N}'(u) + \text{In}(u) + \text{Out}(u)$, and $d'(u) + \text{id}(u) + \text{od}(u) = d$. In particular, $d'(v) = \text{id}(v) = 0$, and $\text{od}(v) = d$. If $u \in V_1 = \text{N}(v)$ then $\text{id}(u) = 1$ and so $d'(u) + 1 + \text{od}(u) = d$, and thus $\text{od}(u) = d - 1 - d'(u) = |\overline{\text{N}}'_1(u)|$.

A pair of distinct vertices $x, y \in V_i$ will be *below* a pair of distinct vertices $a, b \in V_{i+1}$ if $a \in \text{Out}(x)$ and $b \in \text{Out}(y)$. In this case we also say that a, b are *above* x, y . Note that, if $a, b \in V_{i+1}$ is not above some pair $x, y \in V_i$, there must be a unique $z \in V_i$ with $a, b \in \text{Out}(z)$. We will be most interested in the case where $i = 1$ and $ab \notin E_2$.

For other graph-theoretic definitions and concepts not given here, see [24], for example.

1.1 Switches

As described above, an established approach to the generation of graphs with given degrees is to use local edge transformations known as switches. The process is *irreducible* if any graph in the class can be obtained from any other by a sequence of these local transformations. Here we will consider three possibilities for this local transformation.

In a switch, a pair of edges xy, wz of graph $G = (V, E)$ are chosen at random in some fashion, and replaced with the pair xw, yz , provided these are currently non-edges and that the vertices x, y, w, z are distinct. See Figure 1. We make no other assumptions about $G[\{w, x, y, z\}]$. Clearly switches preserve vertex degrees, since each vertex in the switch has one edge deleted and one added, and all other vertices are unaffected.

Taylor [20] proved that the set of graphs with given degrees is connected under switches. (See also [23], [24, Thm. 1.3.33], where switches are called “2-switches”, and [16] for a more constructive proof.) Cooper, Dyer and Greenhill [5] showed rapid mixing of the switch Markov chain for regular graphs, and a generalisation to some (relatively sparse) irregular degree sequences was given in [13]. Switches can easily be restricted to preserve bipartiteness, by requiring that $\{w, y\}$ (or equivalently $\{x, z\}$) belong to the same side of the bipartition. In fact, the first use of switches as the transitions of a Markov chain was for bipartite graphs [15].

If we wish to generate only *connected* graphs, we may use the *flip*. This is defined in the same way as the switch, except that we specify that wy must also be an edge. See Figure 2. Note that a flip is a restricted form of switch which cannot disconnect the graph. Tsuki [23] proved that the flip chain is irreducible for 3-regular connected graphs, while the corresponding result for d -regular graphs was proved by Mahlmann and Schindelbauer [17], for any $d \geq 3$. Cooper, Dyer, Greenhill and Handley [8] subsequently showed rapid mixing for regular graphs




Fig.1 A switch

of even degree. Note that flips are not well-defined on bipartite graphs, since $\{w, y\}$ clearly cannot be on the same side of a bipartition. Mahlmann and Schindelbauer also considered other restricted forms of switches, where there must be a k -edge path between w and y . The flip chain corresponds to $k = 1$, while the “2-flipper” with $k = 2$ preserves connected bipartite graphs.

Irreducibility of the 2-flipper was proved in [17], but the idea does not seem to have been considered subsequently.




Fig.2: A flip

In [6], a different restriction of switches was introduced, designed to ensure that every switch changes the set of triangles in the graph. The definition is as for switches, except that x and w must have a common neighbour, which we denote by v . This is a *triangle switch*, which we abbreviate as Δ -switch. Every Δ -switch makes (creates) or breaks (destroys) at least one triangle. Again, we make no further assumption about $G[\{v, w, x, y, z\}]$. Clearly, Δ -switches do not preserve bipartiteness, since bipartite graphs have no triangles.

Specifically, if the 4-edge path $yxvwz$ is present in the graph and the edges xw, yz are absent, a *make triangle switch* at v , denoted Δ^+ -switch, deletes the edges xy, wz and replaces them with edges xw, yz , forming a triangle on v, x, w . The Δ^+ -switch is illustrated in Figure 3, reading from left to right. Conversely, if the edge yz and the triangle on v, x, w are present in the graph, such that the edges xy, wz are both absent, then a *break triangle switch* at v , denoted Δ^- -switch, deletes the edges xw, yz and replaces them with the edges xy, wz . This destroys the triangle on v, x, w . The Δ^- -switch is illustrated in Figure 3, reading from right to left. Note that a Δ^- -switch reverses a Δ^+ -switch and vice versa. A Δ -switch which involves v and two incident edges, as in Figure 3, will be called a Δ -switch at v . Note that this is equivalent to a switch in the graph




Fig. 3: The Δ^+ -switch and Δ^- -switch triangle switches at v

$H = G[V_1 \cup V_2]$, if the graph is layered from v , and we will use this equivalence in our arguments below.

Now let $\mathcal{M}_{n,d}$ be the graph with $V(\mathcal{M}_{n,d}) = \mathcal{G}_{n,d}$ and $\{G, G'\} \in E(\mathcal{M}_{n,d})$ if and only if G' can be obtained from G by a single Δ -switch. Then a time-homogeneous Markov chain with state space $\mathcal{M}_{n,d}$ will be called a Δ -switch chain if its transition matrix P satisfies $P(G, G') > 0$ if and only if $\{G, G'\} \in E(\mathcal{M}_{n,d})$. That is, $\mathcal{M}_{n,d}$ is the transition graph underlying any Δ -switch chain. Then our main result is the following.

Theorem 1. *Suppose that $d \geq 3$. Then the graph $\mathcal{M}_{n,d}$ is connected. Equivalently, any Δ -switch chain is irreducible on $\mathcal{G}_{n,d}$.*

2 Small regular graphs

Here we show that Δ -switches connect the set of all fragments on a given vertex set. First we give some properties of fragments which are required in our proof.

Lemma 1. *Let G be a d -regular fragment, with $d \geq 3$. Then G is a connected graph with diameter 2.* \square

Remark 1. We claim only connectedness, but fragments have higher connectivity. It is not difficult to prove 2-connectedness. For each d , we have examples with connectivity only $\lfloor d/2 \rfloor + 1$, and we believe this represents the lowest connectivity. However, we make no use of this, so we do not pursue it further here.

For an even integer $d \geq 2$, we construct the graph $T_{d,d,1}$ as follows. Take a copy of $K_{d,d}$ with vertex bipartition (A_d, B_d) , where $A_d = \{a_i : i \in [d]\}$ and $B_d = \{b_i : i \in [d]\}$. Let M be the matching $\{a_i b_i : i \in [d/2]\}$ of size $d/2$ between $A_{d/2}$ and $B_{d/2}$. Form $T_{d,d,1}$ from the copy of $K_{d,d}$ by deleting the edges of M , adding a new vertex v and an edge from v to each a_i and b_i with $i \in [d/2]$. Then $T_{d,d,1}$ is a d -regular tripartite graph with $2d + 1$ vertices and vertex tripartition $\{v\} \cup A_d \cup B_d$.

For example, $T_{2,2,1}$ is a 5-cycle and $T_{4,4,1}$ is shown in Figure 4.

Lemma 2. *Suppose that $G \in \mathcal{G}_{n,d}$ where $d \geq 3$ and $d + 1 < n < 2(d + 1)$. Let ab be an edge of G .*

(i) *If $n < 2d$ then G has a triangle which contains the edge ab .*




Fig. 4: The graph $T_{4,4,1}$

- (ii) If $n = 2d$ then ab is contained in a triangle or an induced C_4 in G . Furthermore, if G is triangle-free then $G \cong K_{d,d}$.
- (iii) If $n = 2d + 1$ then ab is contained in a triangle or an induced C_4 in G . Furthermore, if G is triangle-free then $G \cong T_{d,d,1}$. \square

We use these properties to prove the following.

Lemma 3. *If $d \geq 3$ and $d + 1 < n < 2(d + 1)$ then $\mathcal{M}_{n,d}$ is connected. Equivalently, the Δ -switch chain is irreducible on the set of fragments in $\mathcal{G}_{n,d}$.*

Proof (Sketch). The flip chain [8,17] is irreducible on all d -regular connected graphs. Furthermore, Lemma 1 proves that all fragments are connected graphs. We show that if $d + 1 < n < 2(d + 1)$ then a flip can be performed using at most three Δ -switches. That is, if G is a fragment and G' is obtained from G by a flip, then there is a sequence of at most three Δ -switches which takes G to G' . The lemma then follows immediately.

There are two types of flip, as shown in Figure 5, which we must consider separately. Solid lines are the edges and dashed lines the non-edges involved in the flip which deletes v_1v_2, v_3v_4 and inserts v_1v_3, v_2v_4 , where v_1v_4 is an edge. \square




Fig. 5: The two types of flip

3 Creating a clique component containing a given vertex

The next result, which we prove in this section, is the core of our proof of Theorem 1.

Theorem 2. *Suppose that $d \geq 3$ and $n \geq 2(d + 1)$. Given any $G \in \mathcal{G}_{n,d}$ and any vertex v of G , let $S = N_G[v]$ be the closed neighbourhood of v in G . Then there is a sequence of Δ -switches which ends in a graph G' which has a clique component on the vertex set S .*

Note that $N_{G'}[v] = N_G[v]$, that is, the closed neighbourhood of v is preserved by this process. This property is used in Section 4.

3.1 Proof strategy

Let C be the component of G containing v . We consider $G(C)$ as being layered from v , in a BFS manner, with layer V_{i+1} above layer V_i , and with v in the bottom layer (see Section 1). We prove that, provided $|C| \geq 2(d+1)$, there is a sequence of Δ -switches such that V_1 remains unchanged, but $|E_1|$ increases monotonically. We repeat the following steps to add edges to E_1 until $G[V_0 \cup V_1] \cong K_{d+1}$.

1. If, before any step below, the component C containing v is a fragment (that is, if $d+1 < |C| < 2(d+1)$), use a Δ^- -switch to increase the size of C to at least $2(d+1)$ without removing any edge in E_1 . This can be done in such a way that G_2 now contains at least one non-edge, as we will prove in Lemma 4.
2. While there is a vertex $u \in V_1$ which is not adjacent to any vertex in V_1 (that is, with $d'(u) = 0$) make a Δ^+ -switch to introduce an edge incident with u in G_1 . That this is always possible will be proved in Lemma 5. After repeating this as many times as necessary, every vertex in V_1 will have an incident edge in G_1 . Thus, every vertex $u \in V_1$ with a neighbour in V_2 will have $1 \leq d'(u) \leq d-2$.
3. If $V_2 = \emptyset$ then $G[V_0 \cup V_1] \cong K_{d+1}$, return the current graph as G' . Otherwise, while there is a non-edge ab in G_2 , insert edges into E_1 as follows:
 - (a) Suppose that there is a unique $x \in V_1$ such that $a, b \in V_2$ are in $\text{Out}(x)$ only. Thus $\text{id}(a) = \text{id}(b) = 1$. Use Lemma 6 to make a Δ^+ -switch which replaces edge xb with yb for some $y \in V_1$, $y \neq x$, thus giving a pair x, y below non-edge ab .
 - (b) Suppose that some pair x, y below ab is a non-edge of G_1 . Use a Δ^+ -switch at v to switch xa, yb to xy, ab ; thus increasing the number of edges in G_1 , as in Lemma 7.
 - (c) Now suppose that every pair x, y below ab is an edge xy of G_1 . Choose one such pair and use the Δ -switch at v of Lemma 8 to make xy a non-edge. Then use a Δ -switch at v to switch xa, yb to xy, ab .
4. If $\ell = |V_2| \geq d+1$ then there are necessarily non-edges in G_2 . If $\ell = d$ and $|C| = 2(d+1)$, then $V_3 = \{u\}$, for some u , and $V_2 = N(u)$. Again there are necessarily non-edges in G_2 . In either case go back to Step 3 above.
5. If we reach here, then $\ell \leq d$ and V_2 is a complete graph K_ℓ . If $V_3 = \emptyset$ then $|C| = 1+d+\ell \leq 2d+1$, a contradiction. Thus $V_3 \neq \emptyset$. The case $|C| = 2(d+1)$ was covered in Step 4, so we assume that $|C| > 2(d+1)$. Carry out the steps in Lemmas 9–11 to insert a non-edge into G_2 , and go to Step 3 above.

3.2 Increasing the size of C

Lemma 4. *Suppose that $d \geq 3$ and $n > 2(d+1)$. If vertex v is in a fragment C then there is a Δ^- -switch to increase the size of C to at least $2d+3$ without changing the edges of G_1 . After this switch, G_2 will contain a non-edge. \square*

This procedure does not increase $|E_1|$ but, as we show next, the non-edge in G_2 allows us to increase $|E_1|$ with at most two further Δ -switches. Thus the process outlined in Section 3.1 must terminate in a finite number of steps.

3.3 G_2 has a non-edge

If G_2 has a non-edge, we use the following lemma to increase $|E_1|$.

Lemma 5. *Suppose that $d \geq 3$. Let C be the component of G which contains v , and suppose that $|C| \geq 2(d+1)$. If $u \in V_1$ has $d'(u) = 0$ then we can use at most two Δ -switches to insert an edge in V_1 at u , without altering other edges in E_1 . \square*

Let ab be a non-edge of V_2 above $x, y \in V_1$. We will show that we can rearrange the edges of G_1 as necessary to enable a Δ -switch $axvyb$, replacing xa, yb with xy, ab , inserting an edge xy into E_1 . Lemma 6 deals with the case where ab lies uniquely within $\text{Out}(u)$ for some $u \in V_1$. Lemmas 7 and 8 interchange edges and non-edges in G_1 if necessary. First, we show that we can assume that every non-adjacent pair $a, b \in V_2$ is above some pair in V_1 .

Lemma 6. *Let $d \geq 3$ and $d'(u) \geq 1$ for all $u \in V_1$. Let $a, b \in V_2$ be a pair of distinct non-adjacent vertices such that $\text{In}(a) = \text{In}(b) = \{x\}$ for some $x \in V_1$. Then there is a Δ -switch at v to move b to $\text{Out}(y)$ for some $y \in V_1$, $y \neq x$, without altering E_2 , so that a, b is above x, y .*

Proof. In Figure 6, xb is an edge and so $d'(x) \leq d-2$. Hence there is a non-edge xw for some $w \in V_1$. As $d'(w) \geq 1$ there is some $y \in V_1$ such that wy is an edge. Clearly $y \neq x$. Note that yb is a non-edge because $\text{id}(b) = 1$. (Pairs not shown as an edge or non-edge can be either.) Now switch xb, wy to xw, by . \square




Fig. 6: The switch in Lemma 6, which changes $N(b) \cap V_1$

Lemma 7. *Let ab be a non-edge of G_2 , above a non-edge xy in G_1 . Then there is a Δ^+ -switch to put $xy \in E_1$ without altering any other edges of E_1 .*

Proof. Clearly $axvyb$ is the required Δ^+ -switch. \square

Lemma 8. *Let $d \geq 3$, and let $xy \in E_1$ be such that $\text{od}(y) \geq 1$. Then there is a Δ -switch which makes xy a non-edge, without changing E_2 or decreasing $|E_1|$.*




Fig. 7: The switch in Lemma 7 which inserts xy into E_1 .

Proof. Consider the graph $H = G[V_1 \cup V_2]$. Note that all $u \in V_1$ have $\deg_H(u) = d - 1$. Since $d'(y) = d - 1 - \text{od}(y) \leq d - 2$, but $V_1 \setminus \{y\} = d - 1$, there exists $w \in V_1$ such that $yw \notin E_1$, as in Figure 8. First suppose that $xw \notin E_1$. Let $W = N(w) \setminus \{v\}$ and $X = N(x) \setminus \{v, y\}$, so $|W| = d - 1$, $|X| = d - 2$. Thus there exists $z \in W \setminus X$. So $wz \in E(H)$ and $xz \notin E(H)$, and there is a switch replacing xy, wz by yw, xz . Now $|E_1|$ is unchanged but xy is a non-edge. If $z \in V_1$ then two edges in G_1 are added and two removed. If $z \in V_2$ then one edge of G_1 is added and one is removed. No edges of G_2 are changed, since only z can be in V_2 . Finally, if $xw \in E_1$, let $W = N(w) \setminus \{v, x\}$ and $X = N(x) \setminus \{v, w, y\}$. Then $|W| = d - 2$ and $|X| = d - 3$. The argument now proceeds as above. \square




Fig. 8: The switch in Lemma 8, making xy a non-edge.

If ab was a non-edge above the edge xy , then both $\text{od}(x), \text{od}(y) \geq 1$, so Lemma 8 applies. After performing the Δ -switch from Lemma 8, ab will be above the non-edge xy . Then we use Lemma 7 to re-insert xy and increase $|E_1|$.

3.4 G_2 has no non-edges

If G_2 has no non-edges, we explain how to introduce one without changing E_1 . After this we can apply the results of Section 3.3.

As usual, C denotes the component of G containing v . We assume that $|C| \geq 2(d + 1)$ and that G_2 is complete. If (i) $V_4 \neq \emptyset$ and $\ell \geq 2$ or (ii) all vertices of G_2 have $d - \ell$ edges to V_3 , then we can use Δ -switches to create a non-edge in G_2 . This is proved in Lemmas 10 and 11 respectively. If these conditions are not met, then Lemma 9 describes a procedure which can be repeated until all vertices of V_2 have in-degree one. As a consequence $\ell \geq 2$, and all $u \in V_2$ have $\text{od}(u) = d - (\ell - 1) - 1 = d - \ell$, thus satisfying Lemma 11.

If $\ell \geq d + 1$ then G_2 necessarily has a non-edge, so assume $\ell \leq d$. Then $\text{id}(u) \geq 1$, for any $u \in V_2$, and $d'(u) = \ell - 1$ as G_2 is complete. Thus $\text{od}(u) \leq d - \ell$, so there are at most $\ell(d - \ell)$ edges from V_2 to V_3 . If $\ell = d$ then $V_3 = \emptyset$ and so $|C| = 2d + 1$, a contradiction. Similarly, if $|V_3| = 1$ then $\ell \geq d$ as $V_4 = \emptyset$, and again we have a non-edge in G_2 or a contradiction. If $|C| = 2(d + 1)$ then $V_2 \cup V_3$ is a d -regular subgraph on $d + 1$ vertices, which must be isomorphic to K_{d+1} . But this contradicts the fact that all vertices $u \in V_2$ have $\text{id}(u) \geq 1$. Hence we may assume that $1 \leq \ell \leq d - 1$, $|C| > 2(d + 1)$, and $|V_3| \geq 2$.

Lemma 9. *Suppose that $d \geq 3$ and $|C| > 2(d + 1)$. Further suppose that $1 \leq \ell \leq d - 1$ and $|V_3| \geq 2$, with G_2 complete and $V_4 = \emptyset$. If some $u \in V_2$ has $\text{id}(u) \geq 2$ then there is a Δ^+ -switch which reduces $\text{id}(u)$ by one and moves a vertex of V_3 to V_2 , without altering E_1 . \square*

The above process can be repeated until there is a non-edge in V_2 , in which case we proceed as in Section 3.3, or all vertices on V_2 have in-degree one. In this case $\ell \geq 2$, because there must be at least two vertices in V_1 with out-degree at least one, or else $G[V_1] = K_d$ and we are done.

Thus we may now assume that G_2 is complete with $2 \leq \ell \leq d - 1$, all vertices $u \in V_2$ have $\text{id}(u) = 1$, and $V_3 \neq \emptyset$ (or else C is a fragment). Hence all $u \in V_2$ have $\text{od}(u) = d - 1 - (\ell - 1) = d - \ell$, and there are exactly $\ell(d - \ell)$ edges between V_2 and V_3 . Use the appropriate lemma below, and proceed as in Section 3.3.

Lemma 10. *Suppose that G_2 is complete and $2 \leq \ell \leq d - 1$, all vertices $u \in V_2$ have $\text{id}(u) = 1$, and $V_3 \neq \emptyset$. If $V_4 \neq \emptyset$ then we can apply a Δ^+ -switch to create a non-edge in G_2 without altering E_1 . \square*

Lemma 11. *Let $d \geq 3$ and $|C| \geq 2(d + 1)$. Suppose that G_2 is complete, $2 \leq \ell \leq d - 1$, and all $u \in V_2$ have $\text{id}(u) = 1$. Further suppose that $V_3 \neq \emptyset$ and $V_4 = \emptyset$. Then there is a Δ^- -switch in $G[V_2 \cup V_3]$ which removes an edge of V_2 without altering E_1 . \square*

Remark 2. We can bound the number of Δ -switches required to create a clique component. The steps in Lemmas 5–8 require $\Theta(1)$ Δ -switches for each edge inserted in G_1 , so $\Theta(d^2)$ in total. The steps in Lemmas 9–11 also require $O(1)$ Δ -switches, with the exception of Lemma 9, which could possibly be executed $\Theta(d)$ times between edge insertions in V_1 . Thus the total number of Δ -switches required is $\Omega(d^2)$ and $O(d^3)$. Note that this is independent of n , since at most five layers of G are involved in the process.

4 Relabelling the vertices of clique components

To complete the proof of Theorem 1, we need to show that any graph $X = (V, E_X) \in \mathcal{G}_{n,d}$ can be transformed to any other graph $Y = (V, E_Y) \in \mathcal{G}_{n,d}$ with a sequence of Δ -switches. We will do this by induction on n . It is trivially true for $n = d + 1$, since $\mathcal{G}_{n,d}$ contains only one labelled graph, K_{d+1} . We know from

Lemma 3 that $\mathcal{M}_{n,d}$ is connected for $d + 1 < n < 2(d + 1)$. For $n \geq 2(d + 1)$, we will assume inductively that $\mathcal{M}_{n',d}$ is connected for all $n' < n$.

Choose any $v \in V$. First, suppose that $N_X(v) = N_Y(v)$. We know from Section 3 that we can perform a sequence of Δ -switches to transform X into a graph which is a disjoint union of a clique component on the vertex set $N_X[v]$ and a d -regular graph X' with $n - d - 1$ vertices. Similarly, we can perform a sequence of Δ -switches to transform Y into a disjoint union of a clique component on the vertex set $N_Y[v]$ and a d -regular graph Y' with $n - d - 1$ vertices. Since $N_X(v) = N_Y(v)$, it follows that X' and Y' have the same vertex set. Hence, by induction, there is a sequence of Δ -switches that transforms X' into Y' , as required.

Now suppose that $N_X(v) \neq N_Y(v)$. Using the above procedure, we can assume that $G[N_X[v]] \cong K_{d+1}$, and similarly for Y . We now show how to perform a sequence of switches, starting from X , to ensure that the neighbourhood of v matches $N_Y(v)$.

Let $x \in N_X(v) \setminus N_Y(v)$ and $y \in N_X(v) \setminus N_Y(v)$, where both x and y exist because $|N_X(v)| = |N_Y(v)|$. Since $y \notin N[v]$, it must be the case that y is a vertex of X' . Therefore, let yz be any edge of X' incident on y , and let w be any vertex of $N(v) \setminus \{x\}$, which exists since $d \geq 3$. Note that X has a triangle on the vertices v, w, x , since $G[N_X[v]] \cong K_{d+1}$. Perform the Δ^- -switch shown in Figure 9.




Fig. 9: Swapping x, y in $N_X(v)$

This creates a graph, which we rename as X , such that $N_X(v) \leftarrow N_X(v) \setminus \{x\} \cup \{y\}$. We then perform a sequence of Δ -switches to create a clique component with vertex set $N_X[v]$, using the method of Section 3. After this iteration, again renaming the new graph as X , we find that $N_X[v]$ spans a clique component and $|N_X[v] \cap N_Y[v]|$ has been increased by 1. After at most d repetitions of this process, we have reached a new graph X such that $N_X(v) = N_Y(v)$. We now follow the argument given above for that case, completing the proof.

Remark 3. It might be more efficient to incorporate this step into the procedure described in Section 3. In particular, we could show that x, y can be interchanged as soon as v and x have a common neighbour w . However, as we only need to show that $\mathcal{M}_{n,d}$ is connected, and not that we can find shortest paths in $\mathcal{M}_{n,d}$, we prefer to separate these two steps, for clarity.

Remark 4. We require only one Δ -switch to interchange x, y , but we may have to repeat this d times. Since $O(d^3)$ steps are needed to create a clique component

(see Remark 2), this gives $O(d^4)$ steps in total for each inductive step. This must be repeated in graphs of order $n - i(d+1)$ for $0 \leq i < \lfloor n/(d+1) \rfloor$, that is, $O(n/d)$ times. Thus in total we may need $O(nd^3)$ Δ -switches to connect X with Y .

5 Regular graphs of degree at most two

Theorem 1 excludes the cases $d = 0, 1, 2$. The switch chain is irreducible in all these cases. We will briefly examine the question of connectedness of $\mathcal{M}_{n,d}$ (equivalently, irreducibility of Δ -switch chains on $\mathcal{G}_{n,d}$) when $d = 0, 1, 2$.

If $d = 0$ then the unique graph in $\mathcal{G}_{n,0}$ is a labelled independent set of order n . Hence $\mathcal{M}_{n,0}$ is trivially connected and any Δ -switch chain is trivially irreducible. If $d = 1$ then $G \in \mathcal{G}_{n,1}$ is a matching and n must be even. Now $|\mathcal{G}_{n,1}| > 1$ when $n \geq 4$ is even, but clearly no Δ -switch is possible as no element of $\mathcal{G}_{n,1}$ contains a triangle or a path of four edges. Thus $\mathcal{M}_{n,1}$ is not connected when $n \geq 4$ is even (indeed, $\mathcal{M}_{n,1}$ has no edges in this case).

For $d = 2$, it is not so obvious whether or not $\mathcal{M}_{n,2}$ is connected. We will now examine this case.

If $n < 3$ then $\mathcal{G}_{n,2}$ is empty. For $n \geq 3$, let c_i ($i = 1, 2$) be the number of cycles in $G \in \mathcal{G}_{n,2}$ such that their length modulo 3 is i . Then $n \equiv c_1 + 2c_2 \pmod{3}$. We will say G has *class* (c_1, c_2) .

If $n \geq 3$ then there is at least one class in $\mathcal{G}_{n,2}$. If $n \equiv i \pmod{3}$ ($i \in \{0, 1, 2\}$), then the class $(i, 0)$ exists. Any class is preserved under Δ -switches, since a cycle can either be increased by length 3 by a Δ^+ -switch or decreased by length 3 by a Δ^- -switch. Nothing else is possible. Thus two different classes cannot be connected by Δ -switches. But any graph in the class can be transformed by Δ^- -switches to a “canonical” graph with c_1 4-cycles, c_2 5-cycles and $(n - 4c_1 - 5c_2)/3$ triangles. Note that there are $(k-1)!/2$ distinct labellings of the vertices of a k -cycle.

Lemma 12. *The graph $\mathcal{M}_{n,2}$ is connected if and only if $n \in \{3, 6, 7\}$. Hence any given Δ -switch chain is irreducible on $\mathcal{G}_{n,2}$ if and only if $n \in \{3, 6, 7\}$. \square*

We remark that for $n > 10$ there can be more than two classes in $\mathcal{G}_{n,2}$. For example, $\mathcal{G}_{20,2}$ contains the classes $(5, 0)$ and $(0, 4)$, as well as the two classes $(0, 1)$ and $(2, 0)$ inherited from $\mathcal{G}_{8,2}$.

References

1. Z. Allen-Zhu, A. Bhaskara, S. Lattanzi, V. Mirrokni and L. Orecchia (2018), “Expanders via local edge flips”, in: *Proc. 27th Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA 2016), pp. 259-269.
2. G. Amanatidis and P. Kleer (2020), “Rapid mixing of the switch Markov chain for strongly stable degree sequences”, *Random Struct. Algorithms* **57**, 637–657.
3. M. Bayati, J.H. Kim and A. Saberi (2010), “A sequential algorithm for generating random graphs”, *Algorithmica* **58**, 860–910.

4. B. Bollobás, *Random Graphs* (2nd ed.), Cambridge University Press, Cambridge, 2001.
5. C. Cooper, M. Dyer and C. Greenhill (2007). “Sampling regular graphs and a peer-to-peer network”, *Combinatorics, Probability and Computing* **16**, 557–593.
6. C. Cooper, M. Dyer and C. Greenhill (2019), “Triangle-creation processes on cubic graphs”, [arXiv:1905.04490](https://arxiv.org/abs/1905.04490).
7. C. Cooper, M. Dyer and C. Greenhill (2020), “A triangle process on regular graphs”, [arXiv: 2012.12972](https://arxiv.org/abs/2012.12972).
8. C. Cooper, M. Dyer, C. Greenhill and A. Handley (2019), The flip Markov chain for connected regular graphs, *Discrete Applied Mathematics* **254**, 56–79.
9. P. Erdős and T. Gallai, “Graphs with prescribed degree of vertices (in Hungarian)”, *Matematikai Lapok* **11** (1960), 264–274.
10. T. Feder, A. Guetz, M. Mihail and A. Saberi (2006), “A local switch Markov chain on given degree graphs with application in connectivity of peer-to-peer networks”, in: *Proc. 47th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2006), pp. 69–76.
11. P. Gao and N. Wormald (2017), “Uniform generation of random regular graphs”, *SIAM Journal of Computing* **46**, 1395–1427.
12. S.M. Goodreau, J.A. Kitts and M. Morris, Birds of a feather, or friend of a friend? Using exponential random graph models to investigate adolescent social networks, *Demography* **46** (2009), 103–125.
13. C. Greenhill and M. Sfragara (2018), “The switch Markov chain for sampling irregular graphs and digraphs”, *Theoretical Computer Science* **719**, 1–20.
14. E.M. Jin, M. Girvan and M.E.J. Newman, Structure of growing social networks, *Physical Review E* **64** (2001), 046132.
15. R. Kannan, P. Tetali and S. Vempala (1999), “Simple Markov chain algorithms for generating random bipartite graphs and tournaments”, *Random Structures & Algorithms* **14**, 293–308.
16. C. Lowcay, S. Marsland and C. McCartin (2013), “Constrained switching in graphs: a constructive proof”, in: *2013 International Conference on Signal-Image Technology and Internet-Based Systems*, pp. 599–604.
17. P. Mahlmann and C. Schindelhauer (2005), “Peer-to-peer networks based on random transformations of connected regular undirected graphs”, in: *Proc. 17th Annual ACM Symposium on Parallelism in Algorithms and Architectures* (SPAA 2005), pp. 155–164.
18. I. Miklos, P. Erdős and L. Soukup (2013), “Towards random uniform sampling of bipartite graphs with given degree sequence”, *Electronic Journal of Combinatorics* **20**, #P16.
19. N. Nishimura (2018), “Introduction to reconfiguration”, *Algorithms* **11**, 52.
20. R. Taylor (1981), “Constrained switchings in graphs”, in: *Combinatorial Mathematics VIII*, Springer Lecture Notes in Mathematics **884**, pp. 314–336.
21. K. Tikhomirov and P. Youssef (2020), “Sharp Poincaré and log-Sobolev inequalities for the switch chain on regular bipartite graphs”, [arXiv:2007.02729](https://arxiv.org/abs/2007.02729).
22. A. Tripathi and H. Tyagi (2008), “A simple criterion on degree sequences of graphs”, *Discrete Applied Mathematics* **156**, 3513–3517.
23. T. Tsuki (1996), “Transformations of cubic graphs”, *J. Franklin Inst.* **33(B).4**, pp 565–575.
24. D. West, *Introduction to Graph Theory* (2nd ed.), Prentice Hall, 2000.
25. N. Wormald (1999), “Models of random regular graphs”, in: *Surveys in Combinatorics 1999*, London Math. Soc. Lecture Notes Series **267**, pp. 239–298.